Evidence for the Decay $K^+ \to \pi^+ \nu \bar{\nu}$

About this document

This note is intended to help make the paper reporting E787's new result accessible to an interested audience beyond the high energy physics community. The paper appears in the September 22, 1997 issue of the journal Physical Review Letters.

The note begins with a brief summary followed by an longer section providing the context and motivation for the experiment. The rest of the note is a paragraph-by-paragraph "translation" of our paper, with a section for each paragraph of the original. I have not attempted to add extensive conceptual explanations, but have restricted myself to interpreting jargon, adding detail that was intended to be understood by the physics audience, and "unpacking" the terse style necessitated by Physical Review Letters's length limits. This is, then, not a summary of, but rather an expansion on the paper.

Note: numbered figures are figures from the original paper. Lettered figures have been added for clarity.

Please direct comments on or questions about this document (Is it useful? What might make it more so?) to Peter Meyers (`meyers@viper.princeton.edu`).

What, where, who, when

In a search that began in 1988, Brookhaven Experiment 787 has seen evidence that the positively-charged elementary particle called a "kaon" (K^+, or "kay-plus") can decay into a positively-charged pion (π^+, or "pi-plus"), a neutrino (ν, "nu"), and an antineutrino ($\bar{\nu}$, "nu-bar"). The evidence for this process, written $K^+ \to \pi^+ \nu \bar{\nu}$, consists of a single decay, or "event," that the experimenters have shown is very unlikely to be due to any known process that could mimic $K^+ \to \pi^+ \nu \bar{\nu}$.

The experiment, called E787, is being carried out at the Alternating Gradient Synchrotron (AGS), a proton accelerator at the Brookhaven National Laboratory on Long Island in New York. The accelerator, commissioned in 1960 and extensively improved since then, can deliver a pulse of 6×10^{13} 24 GeV (billion electron volt) protons every 3.5 seconds. The experiment itself is a collaboration of scientists at Brookhaven and Princeton University in the United States, TRIUMF, a laboratory in Vancouver, Canada, and the KEK laboratory and Osaka University in Japan. The original E787 spectrometer was designed and built from 1983 to 1988 and ran from 1988 to 1991. It was rebuilt from 1991 to 1994. The data discussed here were taken in 1995.

Context: flavor-changing neutral currents and the development of the Standard Model

The decay $K^+ \to \pi^+ \nu \bar{\nu}$ has been of interest to high energy physicists since the 1960's, when the absence of decays like it was a major obstacle for the developing theories that would grow into the present "Standard Model," now known to be a wildly successful description of almost all of the interactions of elementary particles. Kaons, like the protons and neutrons that make up atomic nuclei, are made of the more fundamental particles known as quarks. Ordinary matter is made of up just two of the six known quark types (called "flavors"): the up quark and the down quark. The other four quarks, called "strange," "charm," "bottom," and "top," are all heavier, and therefore so are any particles containing them. The kaon is the *lightest* particle to contain one of these *heavier* quarks. A K^+ contains an up quark and an anti-strange quark. Since particles can decay only to lighter particles, for a kaon to decay, the strange quark itself must decay to lighter quarks. Of the four fundamental forces governing the interactions of elementary particles (strong, electromagnetic, weak, and gravitational), only the weak interaction can change the flavor of quarks. Hence, kaons are a natural laboratory for the study of the weak interactions.

Before the mid-1970's, all known weak interactions changed not only the flavor of the interacting particle, but also its electric charge. In the language of particle physics, such a process is called a "charge-changing current," or a "charged-current interaction." For example, 5% of kaons decay as shown in the "Feynman diagram" of Figure A(a), where the K^+ turns into a π^0.

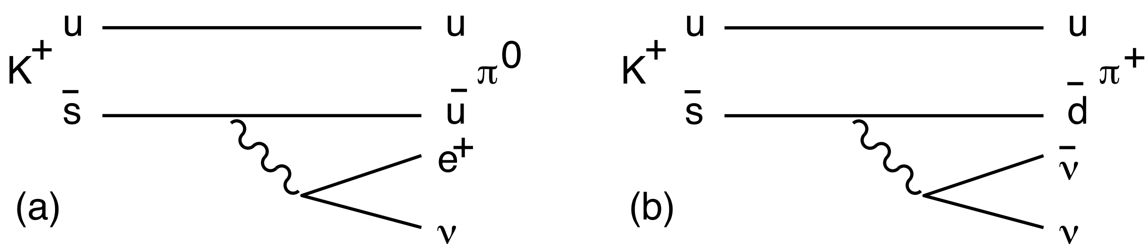

Fig. A. Feynman diagrams for (a) the charged-current process $K^+ \to \pi^0 e^+ \nu$ and (b) the neutral-current process $K^+ \to \pi^+ \nu\bar{\nu}$. Each diagram is read from left to right, and there may be more than one diagram for a given reaction. The wavy lines represent weak interactions. The u quark, a constituent of both the kaon and the pion, is a "spectator" in these reactions, playing no active role.

One of the features of the theories developed to explain the weak interactions was that they required another form of weak interaction, one in which the electric charge of the interacting particle did *not* change. Though the theory predicted that such "neutral-current" interactions would be as common as the known charged-current case, they could have been missed experimentally among the much more common strong and electromagnetic interactions, which are exclusively neutral-current in form. Since there are different quark flavors with the same electric charge, and only weak interactions can change flavor, looking for "flavor-changing neutral currents" seemed the best way to look for the new type of weak interaction. $K^+ \to \pi^+ \nu\bar{\nu}$, where the K^+ changes to a π^+, or, at a more fundamental level, an anti-strange quark ($\bar{s}$, "ess-bar") with charge $\frac{1}{3}$ changes to an anti-down quark ($\bar{d}$, "dee-bar") of the same charge, is such an interaction. (See Figure A(b).)

Searches in the early 1970's for other flavor-changing neutral-current kaon decays failed to find any. Since the searches had the sensitivity to see such decays even if they occurred a hundred-thousand times more rarely than the known charged-current decays, the theory was in trouble. It was rescued by an even more daring theoretical prediction: if a new quark flavor existed, the flavor-changing neutral currents of the new and old quarks would exactly *cancel* each other. With the discovery of this "charm" quark in 1974, the theory became compelling again. Neutral-current weak interactions were found by other means at about the same time.

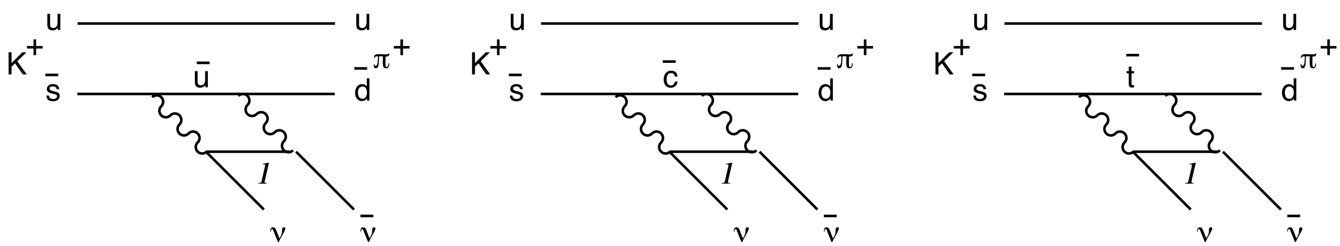

Fig. B. One set of Feynman diagrams for $K^+ \to \pi^+\nu\bar{\nu}$ via "higher-order" weak interactions. There are other sets, each with a diagram containing $\bar{u}$, $\bar{c}$, and $\bar{t}$. In each diagram, l is e, μ, or τ and the neutrinos are the corresponding type.

Where did this leave flavor-changing neutral currents? Though the Standard Model does not allow an anti-strange quark to change directly into an anti-down quark, one could imagine a more complicated ("higher-order") case where the $\bar{s}$ changes into some other flavor via the charged current, and that quark then changes into a $\bar{d}$ via another charged-current interaction. (See Figure B.) It turns out that, while each of the three possible "other" quarks: up, charm, and top, makes an allowed higher-order decay, the *combination* of all three processes again cancels. Unlike the lowest-order case, however, the cancellation here is *not* exact. It would be if the three intermediate quarks had the same mass, but this is not so: the top quark is thousands of times heavier than the up quark. Unlike previous theories of the weak interactions, the Standard Model is able to predict the small-but-not-zero residual rate for $K^+ \to \pi^+\nu\bar{\nu}$. This is usually expressed as the "branching ratio," the fraction of kaons decaying in this mode. In the Standard Model, the branching ratio for $K^+ \to \pi^+\nu\bar{\nu}$ is about 1×10^{-10}, or one in 10 billion, with about a factor of two uncertainty in either direction.

Paragraph 1 — Theoretical interest

Besides testing the bold prediction of a one-in-ten-billion occurrence, the decay $K^+ \to \pi^+\nu\bar{\nu}$ has attracted interest because, within the Standard Model, a large fraction of this small rate is due to a top quark in the transitional role between the $\bar{s}$ and the $\bar{d}$. This means that the rate is sensitive to the strength of the transition (called the "coupling") of top to down quarks. This was recognized even before the top quark was found at Fermilab in 1994. There is a known but poorly-understood hierarchy in the weak interactions among quarks: a quark couples most strongly to the other quark in its own "generation," each generation being a pair of quarks arranged by mass. (See Figure C.)

$$\begin{pmatrix} u \\ d \end{pmatrix} \qquad \begin{pmatrix} c \\ s \end{pmatrix} \qquad \begin{pmatrix} t \\ b \end{pmatrix}$$

Fig. C. Quarks grouped into three generations.

Thus, a top quark couples strongly to a bottom quark, more weakly to a strange quark, and very feebly to a down quark. This makes measurement of the top-to-down coupling, known as V_{td}, very difficult. In fact, it is expected to be impossible at the high-energy accelerators that make and study top quarks directly. This leaves measurement of the $K^+ \to \pi^+\nu\bar{\nu}$ branching ratio as the most direct way to measure the size of V_{td}, a surprising fact, since the energy released in the decay is only a few hundred MeV (million electron volts), where the Fermilab machine that makes top quarks operates at an energy of nearly 2 trillion electron volts per interaction. Other, more indirect, ways of determining V_{td} suffer from the fact that quarks are always bound into strongly-interacting particles (such as kaons and protons). Though the Standard Model contains a fundamental description of strong interactions (Quantum Chromodynamics, or QCD), that theory does not have precise calculational methods for the interactions of low-energy quarks. In $K^+ \to \pi^+\nu\bar{\nu}$, these problems can be avoided by using the measured rate of the decay $K^+ \to \pi^0 e^+\nu$, not a flavor-changing neutral current, but having very similar strong-interaction effects. This would allow a quite clean extraction of the size of V_{td} if the branching ratio were known. Conversely, given less-direct estimates of V_{td}, one can calculate the predicted branching ratio for $K^+ \to \pi^+\nu\bar{\nu}$ given in the previous section.

When E787 began, prior experiments had ruled out branching ratios larger than 1.4×10^{-7}. This meant that there was still a factor of a thousand range in which the Standard Model could fail, either by just being wrong, or because some other process not in the Standard Model occurred. Because the experiment cannot see the neutrinos (which penetrate the spectrometer without interacting at all), other decays of a kaon to a pion and no other detectable particle could look like $K^+ \to \pi^+\nu\bar{\nu}$. The Standard Model has no other such decays, but various proposed extensions to it do. If the undetected part of the event were a single particle, (as opposed to two or more), a large enough sample would allow us to distinguish this case from $K^+ \to \pi^+\nu\bar{\nu}$ by differences in the momentum spectrum of the observed π^+'s.

E787 has published the results of searches using data taken in 1988 and from 1989-91. No $K^+ \to \pi^+\nu\bar{\nu}$ events were observed in that data, allowing us to rule out branching ratios above 2.4×10^{-9}. Since that time, both the kaon beam and the E787 spectrometer have been rebuilt to allow us to increase the sensitivity of the search. In this paper, we report on the analysis of a new data sample taken in 1995, with sensitivity 2.4 times greater than all of our prior data.

Paragraph 2 — Signature and background

The "signature" for $K^+ \to \pi^+\nu\bar{\nu}$ is a K^+ decay to a π^+ of momentum $P < 227$ MeV/c and no other observable product. (The maximum momentum is determined by the mass of the kaon, the mass of the pion, and the fact that neutrinos are massless.) Because we will be working near the limits of our sensitivity, and thus can hope for only a few $K^+ \to \pi^+\nu\bar{\nu}$ events ("signal"), the only way for us to get a convincing observation is to be able to reject events that are *not* $K^+ \to \pi^+\nu\bar{\nu}$ ("background") so that well under one background event survives our analysis. (What does "under one event" mean? If we expect, say, 0.1 event of background, it means that if we had 10 times more data, we would expect one background event.) Estimating the background that would survive our analysis is thus a vital part of the $K^+ \to \pi^+\nu\bar{\nu}$ search.

Major background sources include the copious decays $K^+ \to \mu^+\nu$ (called, in archaic terminology, $K_{\mu 2}$, "kay-mu-two") with a 64% branching ratio and $K^+ \to \pi^+\pi^0$ ($K_{\pi 2}$, "kay-pi-two") with a 21% branching ratio. These are the most common K^+ decays, and they occur billions of times more frequently than $K^+ \to \pi^+\nu\bar{\nu}$ is expected to. The fact that

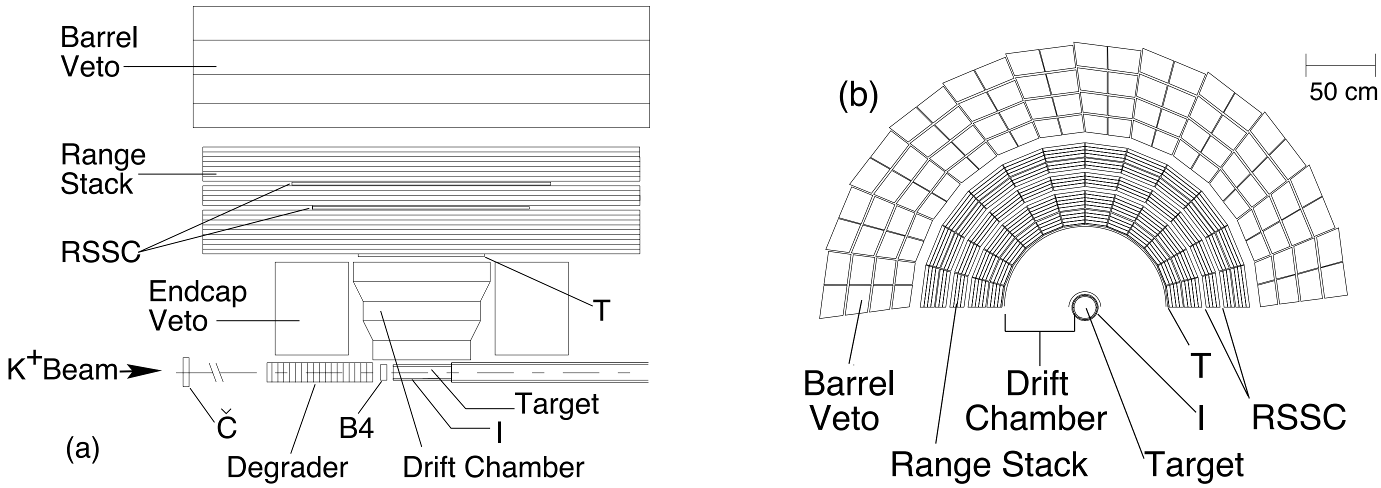

Fig. D. (a) Side and (b) end views of the top half of the E787 spectrometer.

the kaon decays to only two particles in each of these decays is important. The kaons in our experiment are at rest when they decay. In this circumstance, in a "two-body" decay the two resulting particles must have equal and opposite momenta, and that momentum is always the same value. (The spectrum is said to be a "peak".) That value is $P = 236$ MeV/c for $K_{\mu 2}$ and $P = 205$ MeV/c for $K_{\pi 2}$. In contrast, in the three-body decay $K^+ \rightarrow \pi^+\nu\bar{\nu}$, the π^+ has a continuous momentum spectrum extending from zero to 227 MeV/c.

The only other important background sources involve events where what we see isn't a K^+ decay at all. The first occurs when the beam particle entering the detector is a π^+ rather than a K^+. If that π^+ scatters into the spectrometer it can mimic the $K^+ \rightarrow \pi^+\nu\bar{\nu}$ signature. The second occurs when a K^+ undergoes a "charge exchange" (CEX) reaction, becoming a long-lived neutral kaon (K_L^0, "kay-long"), that then decays $K_L^0 \rightarrow \pi^+ l^- \overline{\nu}$, where l is an electron (e^-) or a muon (μ^-).

To suppress the backgrounds, techniques were employed that incorporated redundant kinematic and particle identification measurements and efficient elimination of events with additional particles. Such redundancy, many ways to attack each background, was absolutely necessary for the many orders of magnitude of background rejection needed.

Paragraph 3 — The E787 spectrometer

Figure D shows side and end views of the upper half of the E787 spectrometer. Kaons of momentum 790 MeV/c were delivered to the experiment at a rate of 7 million per 1.6-s pulse of the AGS. The kaon beam line incorporated two stages of purification by particle mass, resulting in a pion contamination of only about 25%. The kaons were detected and identified by Čerenkov (Č in figure D), and energy loss (B4) counters (both particle-type sensitive) and tracking (position-sensitive) counters. The kaons passed through a block of material (Degrader), which slowed them so they stopped in a target of 5-mm-square plastic scintillating fibers. (Scintillator is a material that emits light when penetrated by charged particles. The light is then detected and measured.) About 20% of the incident kaons survived this procedure and decayed at rest in the target. Measurements of the momentum (P), range (R), and kinetic energy (E) of charged particles from the kaon decays were made using the target, a central tracking chamber (Drift Chamber), and a cylindrical "range stack" with 21 layers of plastic scintillator and two layers of tracking chambers (RSSC). The momentum of a particle was measured by tracking it as it curved in the uniform 1-T magnetic field applied parallel to the beam direction throughout the spectrometer. The range was measured using the tracking information and the number of range stack layers penetrated. Kinetic energy was measured by adding up all the light observed from the scintillators along the track.

Pions were distinguished from muons by kinematics (comparing range, momentum, and energy to infer mass) and by observing the $\pi^+ \to \mu^+ \to e^+$ decay sequence in the range stack where the pions came to rest. This was done with "transient digitizers" (TD) that recorded a trace of the light produced in the scintillator, allowing measurement of separate flashes from the π^+ when it enters a scintillator and stops, the μ^+ produced when the π^+ decays [after typically 26 ns (billionths of a second)], and the e^+ produced when the μ^+ decays [after typically 2.2 μs (millionths of a second)].

Photons were detected in a calorimeter (Barrel Veto and Endcap Veto), in which photons interacted and the energy of the resulting particles were measured, completely enclosing the spectrometer. The calorimeter consists of a "barrel" detector made of lead and scintillator sandwiched together and CsI crystal detectors covering each end. In addition, photon detectors were installed along the beam line, including a Pb-glass Čerenkov detector (insensitive to the kaon beam that passes through it) just upstream of the target.

Paragraph 4 — Rejection of background

To be considered as $K^+ \to \pi^+\nu\bar{\nu}$, in each event we required an identified K^+ to stop in the target, followed, after a delay of at least 2 ns, by a single charged-particle track that was unaccompanied by any other decay product or beam particle. The 2-ns delay guaranteed that the kaon actually stopped and that the beam and outgoing tracks were made by different particles. The particle making the outgoing track must have been identified as a π^+ with P, R and E between the $K_{\pi 2}$ and $K_{\mu 2}$ peaks. Approximate versions of these requirements were imposed electronically by a "trigger" that selected events to record on magnetic tape, and later analysis further refined the suppression of backgrounds. To elude rejection, $K_{\mu 2}$ and $K_{\pi 2}$ events would have to have been reconstructed incorrectly, away from their peaks in P, R and E. In addition, any event with a muon would have to have had its track misidentified as a pion — the most effective weapon here was the measurement of the $\pi^+ \to \mu^+ \to e^+$ decay sequence which provided a suppression factor 10^{-5}. Events with photons, such as $K_{\pi 2}$ decays (in which the π^0 decays immediately into two photons), were efficiently eliminated by exploiting the full calorimeter coverage. The chance of missing events with π^0s was only one in a million. A scattered beam pion could have survived the analysis only by misidentification as a K^+ and if the outgoing track were mismeasured as delayed, or if the pion beam track were missed entirely by the beam counters and the pion and scattered into the spectrometer after a valid K^+ stopped in the target. CEX background events could have survived only if the K_L^0 were produced at a low enough energy to remain in the target for at least 2 ns, if there were no visible gap between the beam track and the observed π^+ track, and if the additional e^- or μ^- went unobserved.

Paragraph 5 — Method for estimating surviving background

The data were analyzed with the goal of reducing the total expected background to significantly less than one event in the final sample. In developing the required rejection criteria ("cuts"), we took advantage of the redundant constraints available on each source of background to establish two independent sets of cuts. Applying both sets is intended to leave us with our final sample of only $K^+ \to \pi^+ \nu \bar{\nu}$ events. However, by applying only one set at a time, we can measure the background rejection of that set alone with a sizable sample of events. For example, $K_{\mu 2}$ (including the related decay $K^+ \to \mu^+ \nu_\mu \gamma$) was studied by separately measuring the rejections of the TD particle identification and kinematic cuts. The background from $K_{\pi 2}$ was evaluated by separately measuring the rejections of the photon detection system and kinematic cuts. The background from beam pion scattering was evaluated by separately measuring the rejections of the beam counter and timing cuts. The charge exchange background was handled somewhat differently. Measurements of K^+ charge exchange in the target were performed, which were used as input to a computer simulation ("Monte Carlo"). The estimated CEX background was derived from the simulation. The product of the rejections of the two sets of cuts for each background gave the total rejection, and hence the estimated surviving background. This only works if the two sets of cuts are *uncorrelated*, that is, if they do not preferentially remove the *same* individual events. The sets of cuts were chosen to avoid this, and remaining correlations were investigated for each background source and corrected for if they existed.

There are two important advantages of this method to estimate the background. 1) The use of real data, rather than computer simulation, meant that subtle effects of the instruments and rare pathologies were included in the studies even if we didn't explicitly know of them. 2) The method allowed development of the cuts and estimation of the background levels *prior to* inspection of small samples of $K^+ \to \pi^+ \nu \bar{\nu}$ candidate events. This feature, similar to "blind" studies in biomedical research, helped avoid a biased outcome.

Paragraph 6 — Background estimates

In total, 0.08 ± 0.03 background events were expected in the signal region when the final analysis cuts were applied. These were roughly equally divided among the four background sources. Further confidence in the background estimates and in the measurements of the background distributions near the signal region was provided by extending the method described above to estimate the number of events expected to appear when the cuts were relaxed in predetermined ways so as to allow orders-of-magnitude-higher levels of all background types. Confronting these estimates with measurements from the full $K^+ \to \pi^+ \nu \bar{\nu}$ data, where the two sets of cuts for each background type were relaxed simultaneously, tested the independence of the two sets of cuts. With the background enhanced in this way by a factor of about 20, (expected background 20 times the final analysis) we observed 2 events where 1.6 ± 0.6 were expected, and with a factor of about 150 enhancement we found 15 events where 12 ± 5 were expected. This sample is a good representation of events

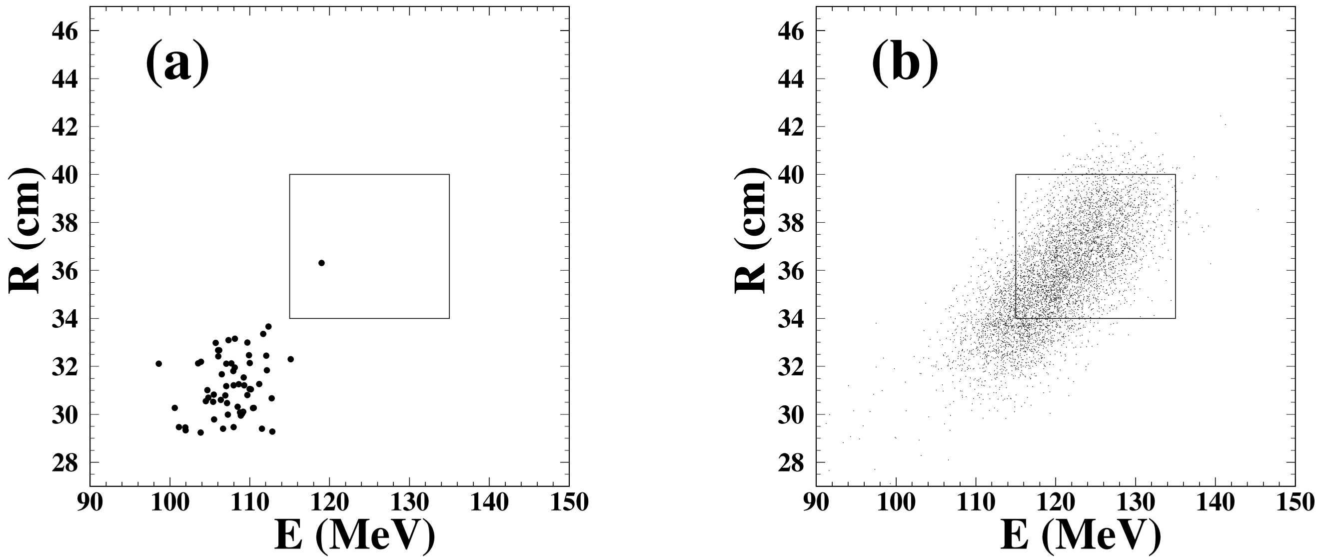

Fig. 1. (a) Range (R) vs. kinetic energy (E) distribution for the $K^+ \to \pi^+ \nu\bar{\nu}$ data set with the final cuts applied. The box enclosing the signal region contains a single candidate event. (b) Computer simulation of $K^+ \to \pi^+ \nu\bar{\nu}$ with the same cuts applied.

that nearly survived the analysis and hence a good place to look for unanticipated sources of background. Under detailed examination, the events were consistent with being due to the known background sources.

Though less than 0.1 event of estimated background was considered adequate for the $K^+ \to \pi^+ \nu\bar{\nu}$ search, we still had additional background rejection capability. Therefore, prior to looking in the signal region, we established several sets of ever-tighter criteria which were designed to be used only to interpret any events that fell into the signal region.

Paragraph 7 — Results of the analysis

Figure 1(a) shows range vs. kinetic energy of the single charged track in events surviving all other analysis cuts, including the requirement that the measured momentum be in the accepted region $211 \leq P \leq 230$ MeV/c. The rectangular box indicates the signal region, specified as range $34 \leq R \leq 40$ cm of scintillator (corresponding to $214 \leq P_\pi \leq 231$ MeV/c for a pion) and energy $115 \leq E \leq 135$ MeV ($213 \leq P_\pi \leq 236$ MeV/c). This box encloses the upper 16.2% of the expected $K^+ \to \pi^+ \nu\bar{\nu}$ spectrum. One event was observed in the signal region. The residual events below the signal region clustered at $E = 108$ MeV were due to $K_{\pi 2}$ decays where both photons from the π^0 had been missed. The number of these events is consistent with estimates of the photon detection inefficiency.

Paragraph 8 — The candidate

A reconstruction of the candidate event is shown in Figure 2. Measured parameters of the event include $P = 219.1 \pm 2.9$ MeV/c, $E = 118.9 \pm 3.9$ MeV, $R = 36.3 \pm 1.4$ cm, and decay times $K \to \pi$, $\pi \to \mu$ and $\mu \to e$ of 23.9 ± 0.5 ns, 27.0 ± 0.5 ns and 3201.1 ± 0.7 ns, respectively. Figure 3 shows the momentum and K^+ decay time of the event. To

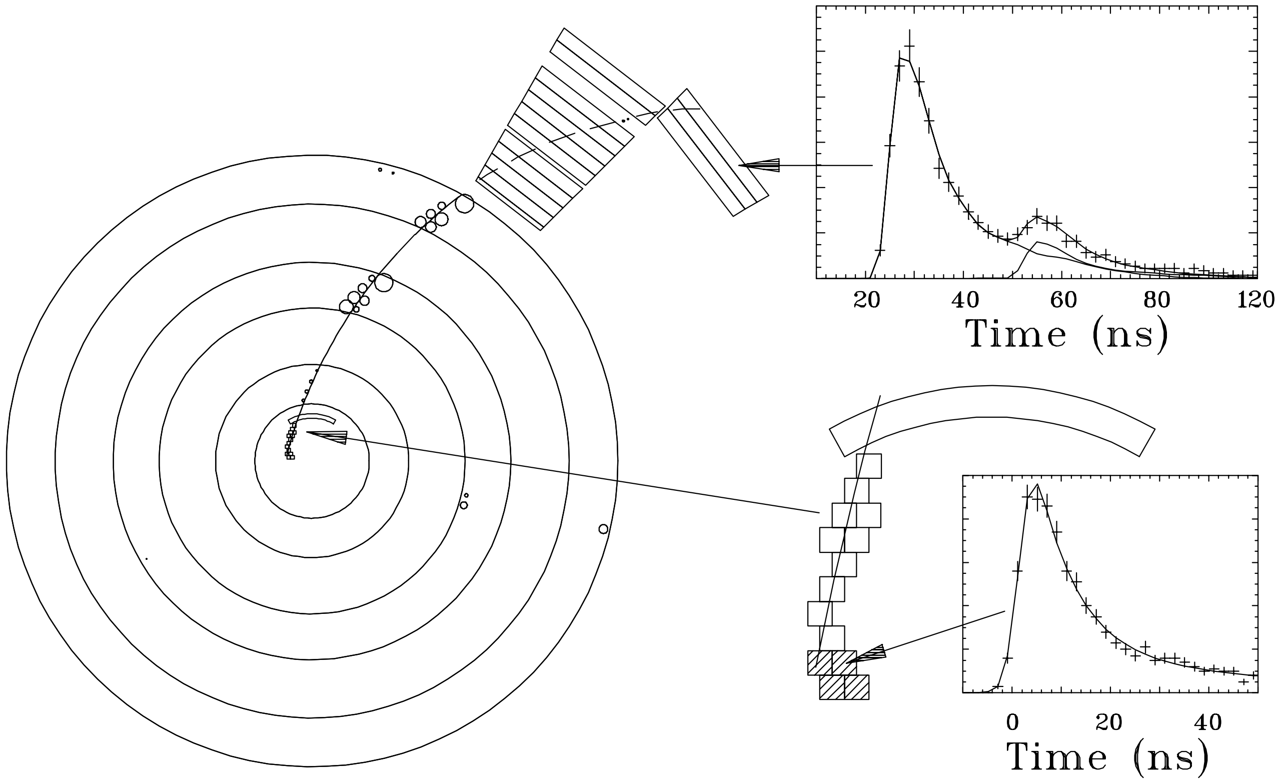

Fig. 2. Reconstruction of the candidate event. On the left is the end view of the detector. In the target and range stack, scintillators with observed energy are outlined. The large concentric circles show the layers of the tracking chamber. The small circles are signals of the passage of a charged particle. The arc is the reconstructed track, which must be tangent to the small circles. At the lower right is a blowup of the target region where the hatched boxes are kaon hits, the open boxes are pion hits, and the inner trigger counter hit is also shown. The light output sampled every 2 ns (crosses) in one of the target fibers hit by the stopping kaon is displayed along with a fit (curve) to the expected pulse shape. At the upper right of the figure is the $\pi^+ \to \mu^+$ decay signal in the range stack scintillator layer where the pion stopped. The crosses are the light output sampled every 2 ns, and the curves are fits for the first (π^+ stopping), second (μ^+ from π^+ decay), and combined pulses.

show where the event lies relative to the major backgrounds, we also show the momentum spectrum when several important cuts are loosened in 3(a) and the time of the outgoing track in events identified as scattered beam pions in 3(b). No significant energy was observed elsewhere in the detector at the same time as the pion. (This might have been evidence of a photon in the event that was barely missed.) The event also satisfied the most demanding criteria designed in advance for candidate evaluation. This put it in a region with an additional background rejection factor of 10. In this region, 0.008 ± 0.005 events would be expected from known background sources while 55% of the $K^+ \to \pi^+\nu\bar{\nu}$ signal is retained. Since, by design, it was unlikely for any background to appear in the signal region, and, by good fortune, it was *very* unlikely (by our pre-selected criteria) for a background event to resemble our candidate, we conclude that we have likely observed

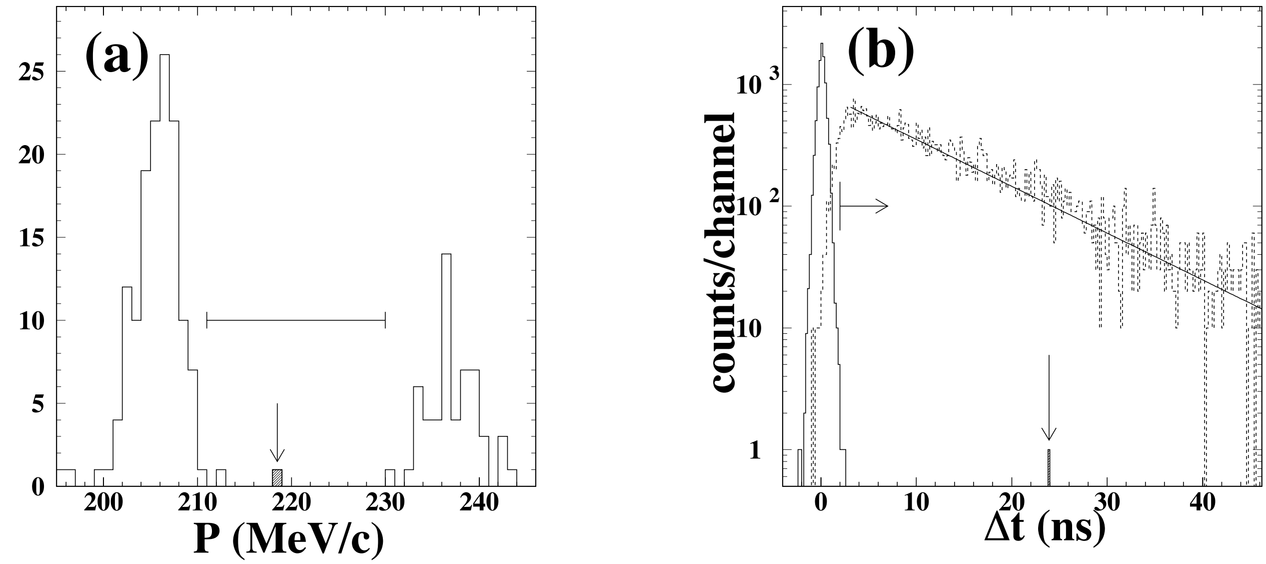

Fig. 3. Event distributions with analysis cuts loosened to show the major backgrounds. (a) The histogram shows the momentum spectrum with backgrounds enhanced by an order of magnitude by loosening the range, photon, and TD particle identification cuts. The peaks are due to $K_{\pi 2}$ and $K_{\mu 2}$. The candidate event (vertical arrow) is shown in relation to the accepted region (horizontal bar). (b) The time difference Δt between the π^+ and the K^+ signals in the target for the candidate event (vertical arrow) and for a sample of events identified in the beam instrumentation as scattered beam pions (solid histogram). The accepted region is $\Delta t > 2$ ns (horizontal arrow). The dotted histogram is the measured time distribution for kaon decays. The straight line shows the known K^+ lifetime.

a kaon decay $K^+ \to \pi^+ \nu\bar{\nu}$.

Paragraph 9 — Calculating the branching ratio

To calculate the branching ratio indicated by this observation, we divide the number of observed $K^+ \to \pi^+ \nu\bar{\nu}$ events (one: our candidate event) by the number of kaon decays observed in total, which is determined by counting the number of kaons entering the target. To this we must apply a correction called the "acceptance," the efficiency for our analysis to find a real $K^+ \to \pi^+ \nu\bar{\nu}$ event. This acceptance is the product of the factors in Table I. For example, we only accept an event as a candidate if the momentum of the pion track is above the $K_{\pi 2}$ peak. This means that we can only see the top 16.2% of the momentum spectrum (the entry in the table listed as "$K^+ \to \pi^+ \nu\bar{\nu}$ phase space"), and thus, for every $K^+ \to \pi^+ \nu\bar{\nu}$ event we see, there were about five others outside the momentum region we looked in. Where possible, we employed calibration data taken simultaneously with the physics data to calculate the factors in Table I, thus automatically taking into account detailed instrumental and time-varying effects. We relied on Monte Carlo simulations only for the fraction of the spectrum accepted, the angular region in which we accepted pions ("Solid angle acceptance"), and fraction of pions lost when they underwent nuclear interactions

Acceptance factors	
K^+ stop efficiency	0.75
K^+ decay after 2 ns	0.813
$K^+ \to \pi^+ \nu\bar{\nu}$ phase space	0.162
Solid angle acceptance	0.386
π^+ nucl. int., decay-in-flight	0.502
Reconstruction efficiency	0.956
Other kinematic constraints	0.713
$\pi - \mu - e$ decay acceptance	0.247
Beam and target analysis	0.659
Accidental loss	0.747
Total acceptance	0.0016

Table I. Acceptance factors used in the measurement of $K^+ \to \pi^+ \nu\bar{\nu}$. The "$K^+$ stop efficiency" is the fraction of kaons entering the target that stopped, and "Other kinematic constraints" includes kinematic particle identification cuts. "Accidental loss" occurs when a particle having nothing to do with an event enters the spectrometer just as the kaon decays.

or decayed before stopping ("π^+ nucl. int., decay-in-flight"). Figure 1(b) shows the R vs. E distribution for many simulated $K^+ \to \pi^+ \nu\bar{\nu}$ events after the final analysis cuts. The acceptance calculation is checked by using our data to make a measurement of the well-known $K_{\pi 2}$ branching ratio. This measurement agrees well with the accepted value and leads us to trust our acceptance calculation to better than 10%.

If the observed candidate event is due to $K^+ \to \pi^+ \nu\bar{\nu}$, the branching ratio is

$$\begin{aligned} B(K^+ \to \pi^+ \nu\bar{\nu}) &= \frac{\text{(Number of observed candidates)/(Acceptance)}}{\text{(Total number of decays)}} \\ &= \frac{1/0.0016}{1.49 \times 10^{12}} = 4.2^{+9.7}_{-3.5} \times 10^{-10}. \end{aligned}$$

The uncertainty attached to the branching ratio is the result of the extremely limited sample of candidate events, that is, one. It means that if many experiments identical to this one were attempted and the true branching ratio were greater than $(4.2+9.7)\times 10^{-10} = 13.9 \times 10^{-10}$, at least 84% of them would have observed more than the one event we saw. (The choice of 84% for the calculation is conventional and corresponds to one "standard deviation"). Similarly, if the true branching ratio were less than $(4.2 - 3.5) \times 10^{-10} = 0.7 \times 10^{-10}$, at least 84% of them would have observed less than one event.

Paragraph 10 — A more exotic possibility?

The likelihood of the candidate event being due to $K^+ \to \pi^+ X^0$, where X^0 is a single, new, massless particle that escapes detection, is small because the π^+ in such a two-body decay would have a momentum of 227.1 MeV/c, far from the observed value. Thus, using the acceptance for $K^+ \to \pi^+ X^0$ of 0.0052 and no observed events in the region $221 < P < 230$ MeV/c, we rule out branching ratios greater than 3.0×10^{-10} for such decays.

Paragraph 11 — An estimate of V_{td}

The fact that the branching ratio for $K^+ \to \pi^+\nu\bar{\nu}$ inferred from our observation overlaps that expected from the Standard Model means that seeing one $K^+ \to \pi^+\nu\bar{\nu}$ event is consistent with Standard Model expectations. Based on this branching ratio, $|V_{td}|$ lies in the range $0.006 < |V_{td}| < 0.06$. E787 has recently collected additional data and the experiment is continuing.